

Degree Examinations 2012-2013

DEPARTMENT OF COMPUTER SCIENCE

Systems Software & Compilers

Paper 3

Time allowed: **90 (ninety) minutes**

Candidates should answer **one** question from **each section**.

Calculators are **not** allowed in this examination.

Do **not** use red ink.

Section A: Compulsory Question

1 (25 marks) SQL

(i) [4 marks] Describe four SQL facilities to implement integrity constraints.

(ii) [21 marks] Consider the insurance database of Figure 1, where the primary keys are underlined.

person (driver-id, name, address)
car (license, model, year)
accident (report-number, date, location)
owns (driver-id, license)
participated (driver-id, license, report-number, damage-amount)

Figure 1: Insurance database schema.

Construct the following SQL queries for this relational database.

- a. List the report numbers of accidents located in York (*2 marks*).
- b. Return the report numbers where the amount of damage is greater than 1000 pounds (*2 marks*).
- c. List the name of the driver, the driver ID, the location, and the amount of damage of each accident (*3 marks*).
- d. Find the total number of people who owned cars that were involved in accidents in 1989 (*4 marks*).
- e. Provide the total number of accidents involving a 'Ford Fiesta' model at each location (*5 marks*).
- f. Find the number of accidents in which the cars belonging to 'John Smith' were involved (*5 marks*).

Section B: Answer ONE question from this section

2 (25 marks) Relational Algebra

catalog (sid, pid, cost)
 supplier (sid, sname, city, country)
 part (pid, colour)

where *sid* stands for "supplier ID"
and *pid* stands for "part Id"

Figure 2: Store catalog schema.

(i) [11 marks] Given the Store catalog schema shown in Figure 2, write the following queries in relational algebra.

- Find the names of suppliers who are based in York, UK (2 marks).
- Find the IDs of suppliers who supply some red or green parts (2 marks).
- Find the IDs of suppliers who supply some red parts and are based in the UK (3 marks).
- Find the IDs of suppliers who supply only red parts (4 marks).

(ii) [5 marks] Given the Store catalog schema shown in Figure 2, for each of the following relational algebra queries, say what they mean:

- $\pi_{sname}(\sigma_{colour='red'}(Part) \bowtie \sigma_{cost < 100}(Catalog) \bowtie Supplier)$ (2 marks).
- $\pi_{sname}(\sigma_{colour='red'}(Part) \bowtie \sigma_{cost < 100}(Catalog) \bowtie Supplier) \cap \pi_{sname}(\sigma_{colour='green'}(Part) \bowtie \sigma_{cost < 100}(Catalog) \bowtie Supplier)$ (3 marks).

(iii) [4 marks] Describe three ways in which inheritance can be represented in a relational data model.

(iv) [5 marks] A relation 'addresses' records information about addresses; for each city and street it gives the corresponding postcode. The intention of addresses is given below:

addresses (city, street, postcode)

and the following dependencies hold:

$$\begin{array}{l} \text{city, street} \rightarrow \text{postcode} \\ \text{postcode} \rightarrow \text{city} \end{array}$$

Based on this, answer the following questions:

- a. Give the candidate key(s) of addresses (*2 marks*).
- b. Is addresses normalised? Assume it is 1NF, but consider the following normal forms: 2NF, 3NF and BCNF. Justify your answer (*3 marks*).

3 (25 marks)

The 'New Home' database maintains information regarding homebuilders and the models that they offer at various locations in specific subdivisions. A subdivision means the division of a lot (that is, parcel of land) into two or more lots. A homebuilder may offer the same model at more than one subdivision, and the price that they offer for a model may be different at each subdivision (think about the difference of prices between a house near a highway compared to a house next to a beach). Each subdivision has lots that are identified by a lot number. Each lot has an associated street address, and size in square feet. When a lot is sold, there is a homebuilder's model associated with the sale. The status of the sale is also recorded.

- The characteristics of a homebuilder include a unique identification number, a name, an address and a phone number. An address is further broken down into a street address, city and postcode. A homebuilder offers several models.
- The characteristics of a subdivision include its name, which is assumed to be unique, the city and the postcode where the subdivision is located.
- The characteristics of a model include an identification number for the model that is unique for its homebuilder. A model from a different homebuilder may have the same identification number. A model also has a name (e.g. 'Three Bedrooms semi detached') and a square footage. A model may be offered at various subdivisions. The price that a model is offered at for a given subdivision is recorded.
- The characteristics of a lot include a lot number that is unique for its subdivision. Each lot has an associated street address, and size in square footage. A sold lot has a particular homebuilder's model associated with it. The status of the sale indicates whether the house construction is 'pending' or 'completed'.

- [12 marks] Provide an ER diagram capturing the conceptual design of the 'New Home' database. You must identify the entities and relations between entities (5 marks). You must justify your design choices (3 marks), identify primary keys (2 marks), and use Chen's notation (2 marks).
- [8 marks] Map your ER design to a relational schema. Justify the choices made to obtain the resulting schema.
- [5 marks] Write the relational algebra expression that returns the list of lots in the 'Bellevue Terraces' subdivision that are available, (i.e. not sold). The returned relation should contain the following attributes (lot number, lot street address, lot size).

DEPARTMENT OF COMPUTER SCIENCE

Degree Examinations 2012-2013

Systems Software & Compilers

Paper 3

Marking Notes

Marking Notes

Answers for Section A: Compulsory Question

Question 1 (25 marks) SQL**Part (i) [4 marks]****[BOOKWORK]** Any of:

- definition of primary and foreign keys
- definition of data domains
- checks of properties involving several attributes of a table
- definition of attributes whose values are unique
- assertions involving attributes of different tables
- triggers

One mark each up to four marks.

Part (ii) [21 marks]**[PROBLEM SOLVING]**

- a. select accident.report-number
from accident
where accident.location = 'York'
- b. select participated.report-number
from participated
where participated.damage-amount > 1000
- c. select person.name, person.driver-id, accident.location,
participated.damage-amount
from participated, person, accident
where participated.driver-id = person.driver-id
and participated.report-number = accident.report-number
- d. select count (distinct owns.driver-id)
from accident, participated, owns
where accident.report-number = participated.report-number
and participated.driver-id = owns.driver-id
and date between date '1989-00-00' and date '1989-12-31'
- e.

Marking Notes

```
select accident.location, count (distinct accident.report-number)
      as number-accident
  from accident, car, participated
 where participated.report-number = accident.report-number
   and participated.license = car.license
   and car.model = 'Ford Fiesta'
  group by accident.location

f. select count (distinct participated.report-number)
  from (participated join owns
        on participated.license = owns.license)
        join person on person.driver-id = owns.driver-id
 where person.name = 'John Smith'
```

Answers for Section B: Answer ONE question from this section

Marking Notes

Question 2 (25 marks) Relational Algebra

Part (i) [11 marks]

- a. $\pi_{sname}(\sigma_{city='York' \wedge country='UK'}(Supplier))$
- b. $\pi_{sid}(\sigma_{colour='red'} \vee \sigma_{colour='green'}(Part) \bowtie Catalog)$
- c. $\pi_{sid}(\sigma_{colour='red'}(Part) \bowtie Catalog) \cap \pi_{sid}(\sigma_{country='UK'}(Supplier))$
- d. $\pi_{sid}(Supplier) / \pi_{sid}(Catalog \bowtie \sigma_{colour \neq 'red'}(Part))$.

Part (ii) [5 marks]

- a. Find the names of suppliers supplying some red part for less than 100 pounds.
- b. Find the names of suppliers such that there is a supplier with that name supplying some red part for less than 100 pounds and a supplier with that name supplying some green part for less than 100 pounds.

Part (iii) [4 marks]

[BOOKWORK] One mark for each of the following:

- by creating a single relation for the superclass, with extra attributes for the subclasses
- by representing inheritance as a one-to-one relationship
- if the inheritance is total and disjoint, by creating a relation for each of the subclasses with all its attributes.

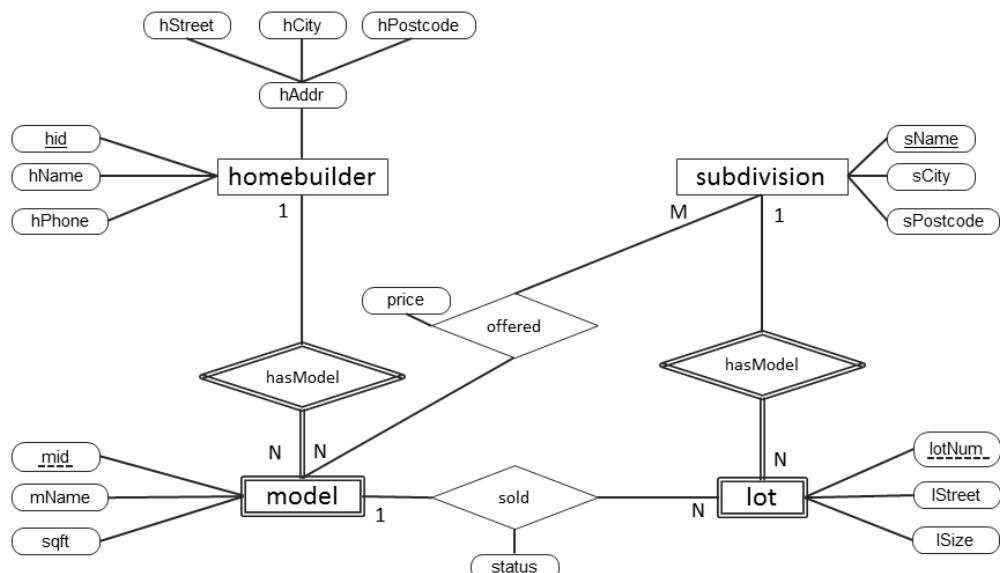
And one mark for the quality of the descriptions.

Part (iv) [5 marks]

- a. (city,street) and (street, postcode)
- b. All attributes are prime so it is 2NF and 3NF. Not in BCNF due to the dependency $postcode \rightarrow city$

Question 3 (25 marks)**Part (i) [12 marks]****[PROBLEM SOLVING]**

The design presented in the diagram below is to be used as a possible answer. Alternative designed with proper justifications of choices should be accepted.

**Part (ii) [8 marks]**

Note this is an indication of what it could be, one of many possible solutions. We should expect six relations, with some differences in attributes. The relation must indicate what are the keys (*2 marks*).

Marking Notes

homebuilder(hid, hName, hStreet, hCity, hPostcode, hPhone) (1 mark).
model(hid, mid, mName, sqft) (1 marks).
subdivision(sName, sCity, sPostcode) (1 mark).
offered(sName, hid, mid, price) (1 mark).
lot(sName, lotnum, lStreet, lSize) (1 marks).
sold(sName, lotnum, hid, mid, status) (1 mark).

Part (iii) [5 marks]

A possible answer is:

```
allTerracesLots  ← πsName,lotnum(σsName='BellevueTerraces'(lot))  
soldTerracesLots ← πsName,lotnum(σsName='BellevueTerraces'(sold))  
availableTerraceLots ← allTerracesLots – soldTerracesLots  
solution        ← πlotnum,lStreet,lSize(availableTerraceLots ⋈ lot)
```